



LAWRENCE  
LIVERMORE  
NATIONAL  
LABORATORY

# **Observations on the Optimality Tolerance in the CAISO 33% RPS Model**

*Yiming Yao, Carol Meyers, Andrea Schmidt,  
Steve Smith, and Fred Streit*

**September 2011**

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

## Observations on the Optimality Tolerance in the CAISO 33% RPS Model

Yiming Yao, Carol Meyers, Andrea Schmidt, Steve Smith, Fred Streitz  
Lawrence Livermore National Laboratory

**Abstract.** In this note we describe our observations on varying the optimality tolerance in the CAISO 33% RPS model. In particular, we observe that changing the optimality tolerance from .05% to .5% leads to solutions over 5 times faster, on average, producing very similar solutions with a negligible difference in overall distance from optimality.

### Background.

In 2008 Governor Schwarzenegger of California issued an executive order requiring that 33 percent of all electricity in the state in the year 2020 should come from renewable resources such as wind, solar, geothermal, biomass, and small hydroelectric facilities. This 33% renewable portfolio standard (RPS) was further codified and signed into law by Governor Brown in 2011.

To assess the market impacts of such a requirement, the California Public Utilities Commission (CPUC) initiated a study to quantify the cost, risk, and timing of achieving a 33% RPS by 2020. The California Independent System Operator (CAISO) was contracted to manage this study. The production simulation model used in this study was developed using the PLEXOS software package, which allows energy planners to optimize long-term system planning decisions under a wide variety of system constraints.

### Optimality Tolerance.

The PLEXOS software works by taking a detailed description of the energy system as its input, in terms of generators, transmission lines, projected loads, reserve requirements, and relevant operational characteristics. It then generates a mathematical model for minimizing the costs of production plus the costs of satisfying the reserves. The mathematical model itself is specified in terms of a *mixed-integer linear program* (MIP), in which the decisions to be made (such as hourly generator unit commitments, over the course of a year) are codified as mathematical variables, and the system constraints are defined in terms of these variables and (static) operational parameters. The costs to be minimized are specified as an objective function defined over these same variables and parameters.

Mixed-integer linear programs are difficult to solve, both in theory and practice, and the time to reach a provably optimal solution can be prohibitively large (weeks or even months), depending on the problem size. The standard industry practice is to establish an *optimality tolerance*, which is a bound within which the objective function is considered acceptably close to optimal for purposes of the study. Note that due to mathematical properties of MIPs, a solution found within the optimality tolerance is in fact provably close (within the tolerance bound) to the *global* optimal solution. The PLEXOS software calls a back-end commercial MIP solver to find these solutions, and post-processes the results.

The PLEXOS model used in the 33% CAISO RPS study is extremely large in scope, covering more than 2000 generators across the entire western WECC grid and associated load, storage, transmission, and

operational requirements over every hour of the simulated year. The mathematical model generated by PLEXOS is correspondingly massive, containing over 225,000 variables and 400,000 constraints. A significant number of these variables (over 34,000) are additionally constrained to being integers, a key fact that strongly influences the running time of the MIP solver. As such, finding a provably optimal solution is not practically feasible and an optimality tolerance must be established.

The optimality tolerance used in the CAISO 33% RPS study is set at .05%, which means that the MIP solver will run until it finds a solution of cost within a provable bound of .05% of optimal. This is a very stringent bound (in some practical applications a tolerance of up to 1% is considered acceptable), and was possibly chosen to be consistent with the tolerance level used in CAISO daily operational models. In addition, the optimization model used to determine each day's unit commitments is constrained to an 8000-second time limit, after which the best solution that is found is returned by the MIP solver, regardless of the specified optimality gap. This was to ensure that daily runs did not take too long.

### Our Findings.

We ran a number of optimization analyses using the PLEXOS model and varying the optimality tolerance level. Because the CAISO 33% RPS analyses are specified in terms of monthly optimization models, we chose five representative months out of the year (January, February, April, August, and November) and averaged the results across months using 6-10 production runs for each month. The monthly models were run in parallel (4-core) on a supercomputer at Lawrence Livermore National Laboratory; there were multiple runs for each month because the underlying MIP solver was run in opportunistic parallel mode, which produces the fastest times but leads to some variability in the actual solutions returned. Our results on solution times are shown in Figure 1.




Figure 1. Average runtime per month, by optimality tolerance level

As can be seen in the figure, the PLEXOS models solved an average of over 5 times faster using a .5% tolerance level instead of a .05% tolerance level. This performance improved only a minimal amount upon further relaxing the tolerance to a 1% level.

Next we investigated the *actual* optimality gap for solutions found at each of these tolerance levels. Because of the 8000-second time limit on daily runs, it was possible for the returned solutions to exceed the optimality tolerance level, if the corresponding runs took too long. Our findings are in Figure 2.




Figure 2. Actual optimality gap obtained, by optimality tolerance level

We observed that at a .05% tolerance level, a highly significant number of daily runs were terminated by hitting the 8000-second time limit, which caused the actual optimality gap to be substantially greater than .05% (three times greater, in fact). For all of the months in the sample, the average optimality gap at a .05% tolerance level was actually greater than .05%. In contrast, at the .5% and 1% tolerance levels, all of the daily runs completed within the 8000-second time limit and most often the solutions found were significantly better (closer to optimality) than the threshold established by the tolerance level.

Finally, we examined the *similarity* of optimal solutions produced under the different tolerance levels. Figure 3 gives an example: it plots the production of generators in the PG&E\_BAY class using generation levels obtained with a .05% tolerance level on the x-axis, and a .5% tolerance level on the y-axis. If the solutions were the same, we would expect a perfect linear correlation for all 44 generator data points.




Figure 3. Comparison of generation levels at PG&E\_BAY in solutions from .5% and .05% tolerance levels

What we observe is actually not far from a perfect correlation (in this case, the correlation coefficient is 0.986, which is extremely close to a perfect value of 1). This shows that the solutions obtained under these different tolerance levels are not substantially different at all. While there are the occasional outliers (one of which can be observed in the chart above), the vast majority of variables are at or near the same values in solutions obtained using the .05% and .5% tolerance levels. We have replicated the same analysis for other variables of interest and the same trend holds true; for the 1% tolerance level, the solutions are also fairly similar in general although not as close as under the .5% tolerance level.

## **Conclusions.**

Our results indicate that changing the optimality tolerance in the CAISO 33% RPS model from .05% to .5% leads to solutions over 5 times faster, on average, producing very similar solutions with a negligible difference in overall distance from optimality. In other words, a greatly significant amount of time could be saved in executing this and future studies using similar models simply by relaxing the optimality tolerance from .05% to .5%. We observe that further increasing the tolerance to 1% does not seem to help nearly as much, both in terms of overall runtime and distance from optimality, and thus would probably be unnecessary.

Lastly, we note that there is actually a substantial amount of uncertainty in the data used in the 33% RPS study; projected load profiles for the year 2020 are inherently highly uncertain. It is possible that the .05% tolerance level was established in running CAISO operational models, which have a great deal less uncertainty, and that this tolerance level might be more natural in such operational models rather than long-term planning studies. This further suggests that using a .5% tolerance level might make intuitive sense in such studies, in addition to the substantial computational savings. As with many aspects of modeling, determining the “right” tolerance level can be an art derived via iteration as much as a science, and we hope our findings add a useful voice to this discussion.